

**VERY HIGH SPEED PAGE OPERATIONS IN  
INDIRECT ACCESSED MEMORY SYSTEMS**

**BACKGROUND OF THE INVENTION**

**Cross Reference to Related Application**

[0001] This application is a continuation-in-part of U.S. Patent Application Serial No. 09/795,414 filed February 28, 2001 which is a continuation-in-part of U.S. Patent Application Serial No. 09/627,516 filed July 28, 2000 which claims the benefit of U.S. Provisional Patent Application Serial No. 60/185,736 filed February 29, 2000.

**Field of the Invention**

[0002] The present invention relates in general to data processing systems and in particular to data processing systems which dynamically allocate main memory for use by the operating system software and application software. Still more particularly, the present invention relates to systems which can directly manipulate pages of memory by modifying the translation table associated with the dynamically allocated main memory.

**Discussion of the Prior Art**

[0003] In computer systems it is customary that there be one-to-one correspondence between the memory address produced by the processor and a specific area in the physical memory of the system. It is an error for the processor to request access to an address which does not have an associated physical memory area. This limits the operating system and applications to an address space determined by the actual physical memory installed in the system. Modern computer systems have overcome this limitation through the use of virtual memory which implements a translation table (TT) to

map program addresses to real memory addresses.

[0004] With virtual memory the program works in an address space limited only by the processor architecture. It is a function of the operating system to ensure that the data and code a program is currently using is in main memory and that the translation table can map the virtual address to the real address correctly. In a virtual memory system the allocation of memory is most commonly performed by the operating system software. This requires an interrupt of the instruction sequence so that the privileged kernel code can allocate physical memory to the area being accessed so that normal program flow can continue without error. This interrupt and the kernel processing to allocate physical memory requires a significant amount of processing time and upsets the normal pipelining of instructions through the CPU.

[0005] There currently exist schemes for reducing operating system process interruptions. For instance, the reference entitled "Design and Analysis of Internal Organizations for Compressed Random Access Memories" by Peter A. Franaszek and John T. Robinson, IBM Research Report RC21146(94535), dated October 28, 1998, describes a low level main memory design for storing compressed data that includes a directory portion and a collection of fixed size blocks which are used to store lines in compressed format. In the memory storage scheme described herein, highly compressible lines may be stored entirely within a directory entry; otherwise, the directory entry points to one or more of the fixed size blocks which are used to store the line in compressed format. The system further makes use of page tables which translate virtual addresses to real addresses which correspond to the location in the directory of the directory entry for the line and which includes information pertaining to blocks holding a compressed line. Specifically, the information in a directory entry includes flags, fragment combining information, and, assuming fixed size entry structure pointers to one or more fixed size blocks. On a cache miss, the memory controller and decompression hardware finds the blocks allocated to store the compressed line and dynamically

decompresses the line to handle the miss. Similarly, when a new or modified line is stored, the blocks currently allocated to the line are made free (if the line currently resides in the RAM), the line is compressed, and then stored in the RAM by allocating the required number of blocks.

[0006] Furthermore, U.S. Patent No. 5,761,536 is directed to a memory organization technique utilizing a compression control device for storing variable length objects (compressed memory) in fixed-size storage blocks by enabling fixed size storage blocks to receive remaining portions (leftover compressed memory pieces or fragments) of variable length objects that take up less than a full fixed-size storage block. The system thus reduces memory fragmentation.

[0007] U.S. Patent No. 5,864,859 is directed to a compression store addressing technique for storing variable length objects (compressed lines, each representing, e.g., 1/4 of a page) in fixed size blocks so that accessing an individual line may be accomplished quickly and with little change to existing software. In particular, the beginning of any line within a page may be accessed with a single pointer plus an offset. Associated with the compression store is a list of free or available blocks (free list) which is accessed for enabling variable length object storage.

[0008] Aforementioned commonly-owned, co-pending U.S. Patent Application No. 09/627,516 addresses the desirability to provide a mechanism that enables the physical memory to be dynamically allocated in a manner such that the interruption in program flow is eliminated, and, furthermore, provides a mechanism for facilitating the management of memory pools so that the various processes and users may share the system resources fairly.

[0009] It is the case that certain simple operations in computer systems such as clear block and move block are commonly performed on large blocks a contiguous memory

bytes called pages. These operations often require that the processor access each individual byte in the page or pages. This makes the operations take an appreciable amount of time and cause the total contents of the page or pages to be loaded into each level of the memory cache hierarchy. This cache damage causes further performance impacts since it evicts other data and programs which are being used by the processor. Subsequent accesses of this purged data will take additional time to reload the appropriate data into the caches

[0010] It would thus be highly desirable to provide in a computer system using dynamically allocated physical memory and a translation table for managing this memory, a mechanism for performing these page operations without requiring the use of the processor.

### **Summary of the Invention**

[0011] It is an object of the invention to provide, in a computer system using dynamically allocated physical memory and a translation table for managing this memory, a mechanism for performing the page operations without requiring intervention of the processor.

[0012] It is another object of the invention to provide, in a computer system using dynamically allocated physical memory and a translation table for managing this memory, a mechanism under the control of hardware for performing the page operations by manipulating only the translation table.

[0013] According to the principles of the invention, there is provided a computing system and method employing a processor device for generating real addresses associated with memory locations of a real memory system for reading and writing of data thereto, the system comprising: a plurality of memory blocks in the real memory system for

storing data, a physical memory storage for storing the pages of data comprising one or more real memory blocks, each real memory block partitioned into one or more sectors, each comprising contiguous bytes of physical memory; a translation table structure in the physical memory storage having entries for associating a real address with sectors of the physical memory, each translation table entry including one or more pointers for pointing to a corresponding sector in its associated real memory block, the table accessed for storing data in one or more allocated sectors for memory read and write operations initiated by the processor; and, a control device for directly manipulating entries in the translation table structure for performing page operations without actually accessing physical memory data contents.

[0014] Thus, according to the present invention, the actual data of the pages involved in the operation are never accessed by the processor and therefore is never required in the memory cache hierarchy, thus eliminating the cache damage normally associated with these block operations. Further the manipulation of the translation table will involve reading and writing a few bytes to perform the operation as opposed to reading and writing the hundreds or thousands of bytes in the pages being manipulated.

[0015] Such a method and apparatus of the invention results in a significant savings in time to perform the actual operation and further represents a smaller performance impact on other activities of the processor. Preferably, the method and apparatus of the invention may be implemented in the design of the compressed memory management systems for server devices, PCs and the like, implementing compression/decompression algorithms.

**Brief Description of the Drawings**

[0016] Further features, aspects and advantages of the apparatus and methods of the present invention will become better understood with regard to the following description, appended claims, and accompanying drawings where:

[0017] Figure 1 depicts generally the data processing system 100 with facilities for dynamic management of physical memory and high speed page operations according to the invention.

[0018] Figure 2 illustrates the continuous running SIU and cSIU monitoring process 450.

[0019] Figure 3 depicts the format for the Class Control Register (CCR) which may be written with a page address and the class which is to be assigned to the page.

[0020] Figure 4 illustrates the process 280 for updating the TT and cSIU's when the CCR is written.

[0021] Figure 5 is a schematic representation of a TT entry 156 in the TT table included in the physical memory.

[0022] Figure 6(a) illustrates the process for writing data into a real address in the indirect accessed memory system.

[0023] Figure 6(b) illustrates the process for reading data from a real address in the indirect accessed memory system.

[0024] Figure 6(c) illustrates the process for deallocating the physical memory

sectors in the indirect accessed memory system.

**[0025]** Figure 7 depicts an example memory state 140 in the indirect accessed memory system prior to a page operation.

**[0026]** Figure 8 illustrates the memory state 142 after the completion of the Empty page operation.

**[0027]** Figure 9 is a flow diagram for the Empty Page operation 200 in the indirect accessed memory system.

**[0028]** Figure 10 illustrates the memory state 144 after the completion of the move page operation.

**[0029]** Figure 11 is a flow diagram for the Page Move operation in the indirect accessed memory system.

**[0030]** Figure 12 is a flow diagram for the Swap Page operation in the indirect accessed memory system.

**[0031]** Figure 13 is a flow diagram for the Copy Page operation in the indirect accessed memory system.

**[0032]** Figure 14 illustrates a flow diagram depicting the Read TT Entries operation in the indirect accessed memory system.

Figure 15 depicts the fields of a CMD command register 176 implemented in the indirect accessed memory system.

Figure 16 is a flow chart depicting the flow required to maintain consistent memory and cache state during the performance of a page operation.

### **Detailed Description of the Preferred Embodiment**

[0033] Figure 1 depicts generally a data processing system 100 with facilities for dynamic management of physical memory. Specifically, in the system of Figure 1, there is provided at least one central processing unit (CPU) 110 which reads and stores data to real address locations in a real memory system 150 and includes L1 and L2 level caches 111, 112 for example. Additionally, interfaced between the CPU and the real memory system is a memory controller device 175 and host bridge 180 operating in conjunction to present real memory images to the CPU, and managing the allocation and manipulation of the physical memory.

[0034] According to the invention, the physical memory 150 includes Random Access Memory (RAM/DRAM) which includes various data structures for managing dynamically allocated physical memory and the data which the CPU reads and writes. These data structures in the RAM include: 1) a mapping table 155 such as a Translation Table (TT) which includes entries comprising pointers to a plurality of sectors 160 where the data is stored in the RAM (i.e., provides a mapping from a CPU generated real memory address to the physical address where the associated data is stored); and, 2) a Free List structure 165 which comprises a list of all of the available, unused, sectors in the physical memory. The plurality of fixed blocks of physical memory (hereinafter "sectors") comprise all of the physical memory, except the areas required for the TT 155 and the Free List 165.

[0035] The RAM 150 is managed by the Memory Controller 175 to carry out the read and store requests from the CPU and any I/O operations via the Host Bridge 180. The Memory Controller 175 includes an interface 185 to the RAM which provides the addressing and controlling signals 181 to the DRAM memory. The interface 185 operates in conjunction with a Memory Management Unit (MMU) 187 that accepts read



and store requests from the CPU 110 and the Host Bridge 180 with associated read addresses. The MMU comprises registers and associated logic 190 for maintaining a count of the number of sectors in DRAM which have been assigned to real addresses (i.e., sectors in use or "SIU"). The MMU additionally includes one or more threshold registers (TR) 195 and comparison logic for continuously comparing the count maintained in the SIU with the TR. When the count maintained by the SIU becomes greater than one or more of the TR(s), an interrupt signal 191 is presented to the CPU 110. Preferably, both the TR(s) and the SIU may be read by the CPU software, however, only the TR(s) may be written by the CPU.

**[0036]** A common task of memory managers such as those found in modern operating systems is to control pools of memory so that the various processes and users share the system resources fairly. To facilitate the management of these memory pools, real memory pages may be grouped into classes and the physical memory usage of each of these classes monitored and managed by the dynamic memory system. To facilitate this classing, each page is provided with a tag which associates it with one of the classes. A typical system may have sixteen (16) classes. For each class, a "Class Sectors In Use" (cSIU) register 192 is provided which always maintains the current count of the sectors that are allocated to pages in that class. In addition to the cSIU registers, there is associated with each cSIU one or more "Class Threshold Registers" (cTR) 197 which may be read and written by the CPU. These cTR's are continuously compared with its associated cSIU 192. If the cSIU is found to have a count larger than the cTR value, the memory controller 175 will alert the management software by interrupting the processor.

**[0037]** Figure 2 is a flow chart depicting the continuous running SIU and cSIU monitoring process 450. As illustrated in Figure 9, the continuous running process implements a step 453 wherein the current count maintained by the SIU is compared to the TR(s). If the current count maintained by the SIU is larger than the count maintained

in the TR, then the CPU is interrupted at step 455. Likewise, at step 458 the current count maintained by each of the cSIU's are compared to their associated cTR(s). If any of these comparisons indicate a cSIU count greater than the associated cTR, then the CPU will be interrupted at step 460.

[0038] Associated with the memory classes there is a mechanism which allows software to control which real memory pages are associated with a given memory class. For example, in the memory controller, a "Class Control Register" (CCR) 198 is provided which may be written with page address and the class that is to be assigned to the page in the manner as depicted in Figure 3. When a page reference and a class are written to the CCR the 'C' field of the TT entries for that page are updated to the new value and the 'cSIU' registers of the old and new classes are update to reflect the change

[0039] Figure 4 is a flow diagram depicting the mechanism 280 for updating the TT entries 156 and SIUs when the CCR 198 is written with page and associated class information. As shown in Figure 4, at step 283, there is performed a first step of reading the translation table entry 156 for the referenced block, and, extracting the old class id ("C") and the sector count for the block. Then, as indicated at step 286, there is performed the step of subtracting the sector count from the associated cSIU register 192 (Figure 1) for that old class C. Then, at step 289, the sector count is updated (added) to the cSIU for the new class written into the register. The corresponding entry 156 in the TT 155 is additionally updated with the new class designation at step 291, and the TT entry 156 is written back to the reference location in the TT 155 at step 293.

[0040] According to the invention, the real memory visible to the CPU and the I/O comprises a set of equal sized blocks of addressable words. For convenience, the size of these blocks is chosen to be an integral power of two (2) of the minimum addressable unit in the memory (byte). For the preferred embodiment, the block size is chosen to be 1024 bytes, with each real address comprising a memory of exactly one of these real memory

blocks. The TT has exactly one entry for each block of real memory and is organized in a regular fashion such that a simple hash of the read address will index in the TT to the associated entry. Such a simple hash may comprise the use of higher order bits as the index.

[0041] Referring back to Figure 1, an area of the physical memory 150 is set aside for the TT 155, the size of which is dependent on the size of the real memory presented to the CPU. In the preferred embodiment, each TT entry is 16 bytes and is associated with a 1024 byte block of real memory. Thus, in this case, the TT would be equal to  $16/1024 = 1/64$  the size of the real memory presented to the CPU 110.

[0042] With more particularity, the Free List structure 165 of Figure 1 comprises a list of pointers to all of the physical memory sectors which are not currently assigned to a real memory block, i.e., an unused physical memory referred to as "free space." That is the free list structure 165 comprises a stack of pointers which point to corresponding sectors in the free space and which stack expands and shrinks during system operation as sectors are assigned to blocks or recovered from blocks. That is, the stack grows deeper as sectors are added to the free space and shallower as free sectors are allocated to real memory blocks. In the preferred embodiment the free list comprises a linked list of free sectors including a "Top Of List" (TOL) pointer that points to the top most entry on the free list. This is the first entry used when fulfilling an assignment of sectors to a real memory block. These sectors are themselves members of the free list so that when they have been emptied of pointers they can also be used by real memory blocks. When a new sector is required for holding pointers to free sectors it is simply allocated from the free list and linked into the chain.

[0043] It is the case that a large portion of physical memory is simply a memory array divided up into sectors which can be used to store the data associated with a real memory block. In the preferred embodiment the sector size is chosen as 256 bytes. Thus

four sectors are required to store all the data associated with a real memory block. In a compressed memory system (such as described in commonly-owned U.S. Patent No. 5,812,817) the compressed representation of a block is of variable length depending on the compressibility of the data. In such a system the physical memory allocated to a real address block need only be sufficient to store the compressed data.

**[0044]** As mentioned with respect to Figure 1, a large portion of physical memory is simply a memory array divided up into sectors which can be used to store the data associated with a real memory block. In the preferred embodiment the sector size is chosen as 256 bytes. Thus, for example, four sectors are required to store all the data associated with a real memory block (1024 bytes). In a compressed memory system (such as described in commonly-owned U.S. Patent No. 5,812,817) the compressed representation of a block is of variable length depending on the compressibility of the data. In such a system the physical memory allocated to a real address block need only be sufficient to store the compressed data. As mentioned, fields in the TT entry for a real address indicate the location and size of the physical memory which is assigned to store the data of the real memory block. Preferably, each TT entry is associated with a 1024 byte real memory block and will hold pointers for up to four sectors assigned to the real memory block. Figure 5 is a schematic representation of a TT entry 156. As shown in Figure 5, each TT entry 156 comprises a set of flag bits 122, and four (4) spaces 125a,...,125d for storing respective sector pointers. One of the flag bits, bit E 130, designates that the associated real memory block is empty and that there are no sectors assigned to that block, i.e., the pointers fields 125a,...,125d are empty. Another field comprising one or more bits C 131 designates the class id that the real memory block is associated with. It should be understood that, as part of an initialization routine, a mechanism is provided such that there is no physical memory block associated with the index, i.e., E field, and, further that the C bit field 131 is initialized to a default value. As mentioned with respect to Figure 4, when a page reference and a class are written to the

CCR 198, the "C" field of the TT entries for that page are updated to the new value and the "cSIU" registers of the old and new classes are updated to reflect the change.

[0045] According to the invention, three operations are defined for dynamically allocating physical memory: 1) write data to a real address; 2) read data from a real address; and, 3) empty real block referenced by a real address.

[0046] The process 300 for writing data into a real address is described in accordance with Figure 6(a). As shown in Figure 6(a), the process includes a first step 303 of applying a hash function to a CPU generated real address for determining an index into the TT table 155. Then, at step 306, the TT entry associated with the referenced real memory block is read. Furthermore, at step 306, the associated cSIU register is located as referenced by the C flag in the TT entry. At step 310, a determination is made as to whether the block is empty, i.e., whether the E bit flag 130 indicating no physical memory allocated to the real memory block. If the block is empty, then at step 313, four sectors are obtained from the free list 165 and corresponding pointers are placed in the respective store fields 125a,...,125d of the TT entry. The SIU register 190 and cSIU register 197 in the memory controller are then each incremented by four (4) at step 316 and the E bit flag 130 is cleared in the TT entry at step 320. After step 320, or if, at step 310, it is determined that the block is not empty, the process continues to step 325 where an offset of the real address in the real memory block is calculated. Then, at step 330, the calculated offset is used to index into allocated sector and the requested data is stored in word(s). Thus, from the write data process depicted in Figure 6(a), it is understood that if the block is marked as empty, i.e., there is no physical memory allocated to the block, sectors are assigned to the block before the write is performed. Otherwise, the write is performed to the appropriate word(s) in one of the sectors assigned to the block.

[0047] The process for reading data from a real address is described in accordance with Figure 6(b). As shown in Figure 6(b), the process 350 for reading data from a real

address includes a first step 353 of applying a hash function to a CPU generated real address for determining an index into the TT table 155. Then, at step 356, the TT entry associated with the referenced real memory block is read. At step 360, a determination is made as to whether the block is empty, i.e., whether the E bit flag 130 indicating no physical memory allocated to the real memory block. If the block is empty, then at step 363, a default data value is returned and the process terminates. If, at step 360, it is determined that the block is not empty, then the process continues to step 365 where an offset of the real address in the real memory block is calculated. Then, at step 370, the calculated offset is used as an index into allocated sector and the requested data word(s) is read and returned. Thus, from the read data process depicted in Figure 6(b), it is understood that the data is retrieved from the appropriate place in one of the sectors assigned to the block unless the block is empty. If the block is empty a default value is returned for the read. Note that a read of an empty block could optionally signal a memory access fault to the CPU.

**[0048]** When a real memory block is no longer needed, the operating system may deallocate the physical memory sectors associated with that real memory block by issuance of an "empty real block" command from the processor. The process for deallocating the physical memory sectors that are no longer needed is described in accordance with Figure 6(c). As shown in Figure 6(c), the process 400 for deallocating the physical memory sectors includes a first step 403 of applying a hash function to a CPU generated real address for determining an index into the TT table 155. Then, at step 406, the TT entry associated with the referenced real memory block is read.

Furthermore, at step 406, the associated cSIU register is located as referenced by the C flag in the TT entry. At step 410, a determination is made as to whether the block is empty indicating no physical memory allocated to the real memory block. If the block is empty, then the process terminates. Otherwise, if it is determined that the block is not empty, then at step 413, the four sectors pointed to by the corresponding TT entry are added to the free list 165. The SIU register 190 and the cSIU register 197 in the memory

controller are then each decremented by four (4) at step 416 and the E bit flag 130 is set in the TT entry at step 420. After step 420, the process terminates. Thus, the deallocating physical memory sectors process depicted in Figure 6(c) enables conservation of physical memory by returning the physical memory sectors to the free list for reuse with other real memory blocks.

## **Page Operations**

[0049] In an indirectly addressed memory system as described above, certain common operations may be performed by directly manipulating the entries in the translation table (TT), without actually accessing the data bytes. Block moves and swaps, and clear block are examples of the type of operations which can be performed using this method. The data can be moved by moving only the sector pointers from the TT entry of the source block to the TT entry of the destination block. Additional block operations such the "empty real block" command described above can be implemented using this mechanism. Block control flags may also be manipulated with page operations.

[0050] These types of operations are commonly performed on pages of memory. A page is the unit of allocation used by the virtual memory manager in the operating system to assign real memory resources to processes in the system. Virtual memory pages are usually a fixed size of a few thousand bytes, 4096 bytes is the most often used page size. The block size for dynamic memory allocation is most likely somewhat smaller, e.g., 1024 bytes in an example embodiment. Thus, a page consists of a few blocks, four in the example embodiment described herein. A page operation will therefore be effected on all of the blocks making up a page as a single operation.

[0051] Figure 7 depicts an example memory state 140 in the indirect accessed memory system and will be used as the starting point for several examples of page

operations. Specifically, as shown in Figure 7, there are two allocated memory pages "A" and "B" shown in the TT table 155. As depicted, each of the TT table entries 156 for pages A and B comprise pointers that point to sectors 162 in the available physical memory of the system. The free list 165 additionally includes entries including pointers to free sectors 161 in the memory. An example first page operation is an "Empty Page" operation performed on Page A whereby the page operation will cause the release of sectors 162 in physical memory associated with that page and return them to the Free List 165. Figure 8 illustrates the memory state 142 after the completion of the Empty page operation. As shown in Figure 8, Page A no longer has any sectors associated with it and the sectors which were associated with it prior to the Empty Page operation are now sectors 161' which are pointed to by the Free List 165.

**[0052]** Figure 9 is a flow diagram for the Empty Page operation 200. As shown in Figure 9, at a first step 202, the page to be emptied is first invalidated in all cache levels. Then, at step 204, the corresponding TT entry for a block comprising the page is read and, at step 206 a determination is made as to whether any sectors are attached to the page. If there are sectors attached to the page, then at step 209, these sectors pointed to by the TT are placed on the free list, and the corresponding SIU register and class counters are decremented. Once it is determined that no remaining sectors are attached to the page, then the corresponding TT entry is cleared of pointers at step 211 and the cleared entry is written back to the TT at step 212. As a page may comprise several blocks it should be understood that process steps 204-212 would repeat for each block in the page. It should be understood that for all of the page operations described herein, processes performed upon sectors associated with blocks will be repeated for each block of the page.

**[0053]** An example second page operation is a "Move Page" operation which is performed to move the contents of Page B to Page A. Figure 10 illustrates the memory



state 144 after the Move page operation is completed. As shown in Figure 10, the TT table entries 156 having pointers which were associated with Page A have been released, and moved to the Free List 165 in view of the empty page operation. Then, as depicted in Figure 10, the TT entries for Page A have been modified to reflect the original state of Page B and the TT entries for Page B have been set to the empty state.

[0054] Figure 11 is a flow diagram for the Page Move operation 220 wherein the sectors from the destination page, e.g., Page A, are returned to the free list and the sectors of the source page, e.g., Page B, are moved to the destination page. It is understood in such page operation that the SIU, and class counters are updated appropriately. More particularly, as shown in Figure 11, the destination page (e.g., page A) in all caches are first invalidated at step 222, and, the source page (e.g., page B) is flushed from all caches at step 225. Then, at step 228, all source page and destination page TT entries are read. At step 231, the number of sectors in the destination page is subtracted from the SIU and at step 233, the number of sectors in the source page is subtracted from the source class count. Next, at step 236 the difference between the number of sectors in the source page and the destination page is added to the destination class count. At step 239, a determination is made as to whether any sectors remain attached to the destination page. If sectors remain attached to the destination page, then the sectors pointed to by the destination TT are placed on the free list as indicated at step 241. Otherwise, and until no further sectors remain attached to the destination page, the process continues to step 243 where the source page TT pointers are moved to the same location in the destination TT entry. Then, at step 245, the source TT pointer entries are cleared and the source and destination entries are written back to the TT at step 247. As a page may comprise several blocks it should be understood that process steps 228-247 would repeat for each block in the page.

[0055] Two more page operations are now described that relate to "Swap Page" and "Copy Page".

**[0056]** Figure 12 is a flow diagram for the Swap Page operation 250 wherein the sectors associated with a source page, e.g., Page A, are swapped with sectors associated with a destination page, e.g., Page B. As shown in Figure 12, the source and destination pages are first flushed in all caches at step 252, and, the source page and destination page TT entries are read at step 255. At step 258, the difference between the number of sectors on the source page and the number of sectors on the destination page is calculated. Then, at step 260, the calculated difference from previous step 258 is subtracted from the source page class count, and at step 262, the calculated difference from previous step 258 is added from the destination page class count. Then, at step 264, the pointers between the source TT entry and the destination TT entry are swapped, and at step 265, the corresponding source and destination entries are written back to the TT. As a page may comprise several blocks it should be understood that process steps 255-265 would repeat for each block in the page.

**[0057]** Figure 13 is a flow diagram for the Copy Page operation 270 wherein the sectors from the source page, e.g., Page A, are copied to sectors associated with a destination page. More particularly, as shown in Figure 13, the destination page in all caches are first invalidated at step 272, and, the source page is flushed from all caches at step 275. Then, at step 278, all source page and destination page TT entries are read and, at step 280, the difference between the number of sectors in the source page and the number of sectors in the destination page is calculated. Next, at step 282 the difference calculated in the previous step 280 is added to the SIU and, at step 284, the difference calculated in the previous step 280 is added to the destination class count. Next, at step 285, a determination is made as to whether any sectors remain attached to the destination page. If sectors remain attached to the destination page, then the sectors pointed to by the destination TT are placed on the free list as indicated at step 287. Otherwise, and until no further sectors remain attached to the destination page, the process continues to step 290 where a determination is made as to whether any sectors remain attached to the source

page. If sectors remain attached to the source page, then at step 292, new sectors are allocated to the destination page from the free list, and, at step 294, the sector contents from the source page sectors are copied to the destination page sectors. Otherwise, and until no further sectors remain attached to the source page, the source and destination entries are written back to the TT at step 295. As a page may comprise several blocks it should be understood that process steps 278-295 would repeat for each block in the page.

[0058] From the foregoing flow diagrams of Figures 9, 11, 12 and 13 and the accompanying description, it is understood that the state of the translation table TT includes information important for the efficient usage and management of dynamically allocated memory in the indirect memory access system. In the indirect memory access system 100 of Figure 1, the contents of the TT is not directly available to programs executing on the processor. A page operation is provided that makes this information readily available to such operating system and management software. For example, as shown in Figure 1, the memory controller 175 is provided with a TT Buffer 177 which is a small memory array that is large enough to hold one or more TT entries. The current contents of the TT Buffer may be read by the processor using normal memory access instructions. In addition the TT Buffer 177 may be loaded with TT entries copied from the TT 155 under control of a special page operation command referred to herein as a "Read TT Entries" page operation. In this command, a first register 'AR1' 178, shown in Figure 1, is loaded with the real address of the TT entries of interest and the command to read the TT is issued through a command (CMD) register 176. Once the TT entries are in, a TT Buffer program running in the processor 110 may analyze the entries for information like the number of sectors assigned to a page, or the number of pages which are empty, etc.

[0059] Figure 14 illustrates a flow diagram depicting the "Read TT Entries" operation 470. It is understood that this operation retrieves an instantaneous snapshot of the entries for the requested page and does not affect the state of the memory. As shown in Figure

14, the Read TT Entries operation comprises a first step 472 where the contents of register AR1 are used to index into the TT, and, a second step 474, where the TT entries at the index point are read into the TT Buffer element 177 (Figure 1).

**[0060]** Page operations are performed by directly interacting with the memory management function in the memory controller 175. Three additional control interfaces are provided to the memory controller 175 for this capability: a first Address register (AR1) 178 in the system block diagram 100 of Figure 1 which is loaded with the page address of the memory block being manipulated or the source address for the data being moved or copied; a second Address register (AR2) 179 in the system block diagram of Figure 1 which is used for the destination address for data being moved or copied; and, a command register CMD 176 shown in Figure 1, which enables initiation of the page operation.

**[0061]** Figure 15 depicts the internal fields of the CMD command register 176 including a function code (FC) 167 and a start flag (SF) bit field 168. Initiation of a page operation is enabled after receiving a function code (FC) 167 and setting the start flag (SF) bit 168 in the command register 176. Completion of a page operation is indicated when the memory controller 175 resets the start bit 168 in the command CMD register 176 to indicate that the operation is done.

**[0062]** It should be further understood, however, that although the contents of the physical sectors in memory 160 are not changed during any of the page operations, they do change the contents of real memory as observed by the processors in the system. For example, when a page operation is performed, any caches which maintain copies of any portions of the associated blocks must be updated appropriately. This is performed using the normal cache coherency protocols of the system with coherency actions initiated by the page operation control logic. Certain page operations, "Empty page," for example, modify or destroy the data in real memory as observed by the processors in the system. It

should further be noted that the actual data stored in the physical memory sectors 161 are not modified, instead these sectors are disassociated from the real memory as observed by the processor. When all of the data in a block is destroyed in this manner, the underlying physical memory assigned to the block may be released for usage by another block in the system. The page operation logic interacts with the memory management logic to cause the release of the associated physical memory when the data of that block is destroyed as a result of the page operation, but the actual contents of the physical memory sectors 161 are not changed during this process. Data modifying page operations must also update or invalidate any cached copies of these blocks.

**[0063]** The actual completion of page operations may be deferred until the data in the affected blocks is actually needed. To allow this deferral, any block with a pending page operation is marked in the CTT with an indication that an operation is pending. If another memory operation is initiated to a block with an outstanding page operation, the new operation is deferred until the page operation is completed. Coherency operations associated with a page operation which is pending, must be completed before the page operation completion is indicated in the control register.

### **Coherency Management in Page Operations**

**[0064]** As mentioned, some page operations change the state of the memory system. As another example, the swap page operation changes the state of both the source and the destination pages. In the system depicted in Figure 1 where portions of these pages may be cached in the processor, these caches 111, 112 must be updated in a coherent manner to assure that memory state is maintained consistent in all levels of storage. The flow diagrams depicted in Figures 9, 11, 12 and 13 describe operations which modify memory contents and require coherency actions on the source and destination pages to maintain full coherency. In addition to the cache flush and invalidate operations, accesses to the pages involved in the page operations are blocked until the page operation has been

completed. Figure 16 is a flow chart depicting the flow 475 required to maintain consistent memory and cache state during the performance of a page operation. New accesses to the affected pages are held pending until the page operation is completed. Any ongoing accesses to the affected pages are allowed to complete. The operation required coherency operations are performed, the actual page operation is performed, and, finally any pending accesses to the pages are allowed to proceed.

**[0065]** More particularly, as shown in Figure 16, there is provided a first step 477 which requires blocking of any new accesses to the involved pages and keep the accesses pending until actions are completed. Then at step 479 a decision is made as to whether there has been any active action(s) on the effected pages. If yes, the process waits pending until actions are completed. Once they are completed, the process proceeds to step 482 where the required coherency actions are performed. Then, at step 484, data modifications and data movements required by the page operation are performed. Finally, at step 486, the pending accesses are allowed to complete and the involved pages are unblocked.

**[0066]** In certain cases the coherency of the pages involved in the page operation is known before the operation is started. In these instances, time and loading on the coherency management circuitry may be saved by only performing the coherency operations required. An example of when this is true is when page operations are used to move data between active pages and pages which are used only to store inactive data such as is done with paging stores in virtual memory systems. In this case, the coherency of the inactive pages is known and the relevant actions may be omitted. Table 1 describes the coherency options for each of the Empty, Move, Swap and Copy page operations which change memory state.

Page Operation	Source AR1	Destination AR2
-------------------	---------------	-----------------

		<b>Full Coherency</b>	<b>Partial Coherency</b>
Empty Page	Invalidate	N/A	N/A
Move Page	Flush	Invalidate	none
Swap Page	Flush	Flush	none
Copy Page	Flush	Invalidate	none

**Table 1: Coherency options for Page Operations**

[0067] As shown in Table 1, full coherency is always performed on the source page, but coherency operations on the destination page may be omitted. Referring back to Figure 16, the CMD register 176 is provided with a Destination Coherent (DC) bit 169 that is set when coherency is desired on the destination page. Otherwise, only partial coherency is assured and the coherency actions on the destination page will be omitted.

[0068] While the invention has been particularly shown and described with respect to illustrative and preformed embodiments thereof, it will be understood by those skilled in the art that the foregoing and other changes in form and details may be made therein without departing from the spirit and scope of the invention which should be limited only by the scope of the appended claims.